SKAO

Dress rehearsal review

SQUARE KILOMETRE ARRAY

Exploring the Universe with the world's largest radio telescope

Rohini Joshi

SRC Scientist

SKAO tests

https://github.com/ESCAPE-WP2/rucio-analysis/tree/dress-rehearsal-1

Dress Rehearsal

- **Pulsar observations**
  Simulate two sets of SKA pulsar observations where pulsar search and single pulse search observations are occurring commensally. The data (50MB) is generated at a different cadence and managed via separate scopes.
- **QoS lifecycle**
  This test attempts to replicate a data lifecycle hourly based on QOS labels on RSEs in the datalake.
- **SDC1 data movement**
  Move SDC1 data to more performant storage (FAST QoS) to simulate “making data available” to the end user.

Post dress rehearsal

- **High throughput test**
  Initiate a transfer of a significant volume of data in an attempt to flood the corresponding network link, and understand the achievable throughput between sites.
- **Million file test**
  Intended to be a Rucio stress test

Pulsar Injection tests: Description

- PSS
  - Hourly uploads of 3500 50 MB files, 350 datasets with 10 files.
    In reality, the throughput would be much greater but one batch of 3500 files takes ~50 min so the frequency was set to be hourly.
  - Lifetime of 2 weeks.
  - All datasets were attached to the container SKA_SKAO_TEAM_PSS:ps_trial_17-11-20.
  - **Missed transfers when the auth server was restarted (~15:15 UTC)**
- SPS
  - Hourly uploads of 20 50 MB files with each group of 20 files contained in 5 datasets
  - Lifetime of 1 hour.
  - All files and datasets for this test are grouped under the container SKA_SKAO_TEAM_SPS:sps_trial_17-11-20.
  - Since each run of the test completes in less than 1 minute and the cron offset for this test was 5 past the hour, this test was unaffected by the Rucio auth server restart.

Pulsar tests: Results

PSS:

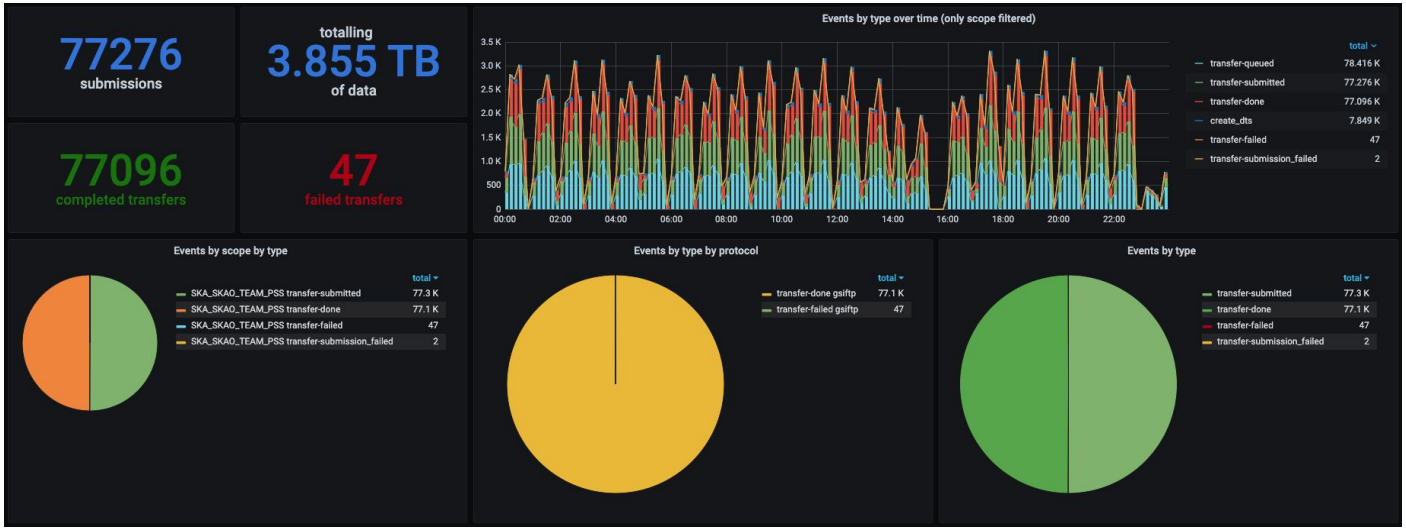

SPS:

SKA Tests

Dress Rehearsal

- Pulsar observations
  Simulate two sets of SKA pulsar observations where pulsar search and single pulse search observations are occurring commensally. The data (50MB) is generated at a different cadence and managed via separate scopes.
- QoS lifecycle
  This test attempts to replicate a data lifecycle hourly based on QOS labels on RSEs in the datalake.
- SDC1 data movement
  Move SDC1 data to more performant storage (FAST QoS) to simulate “making data available” to the end user.

Post dress rehearsal

- High throughput test
  Initiate a transfer of a significant volume of data in an attempt to flood the corresponding network link, and understand the achievable throughput between sites.
- Million file test

QoS lifecycle test

- Upload and replicate based on QoS with a 100MB test file every hour
  - Upload to FAST QoS with a lifetime of 0.5 week
  - Replicate to CHEAP-ANALYSIS QoS with lifetime of 1 week,
  - OPPORTUNISTIC QoS with lifetime of 1.5 weeks, and
  - SAFE QoS with lifetime of 2 weeks.
- All files from this test are grouped in the collection SKA_SKAO_TEAM_QOS24:17-11-2020.
- Config file used here
- Not impacted by issues at GSI-ROOT
- **Rule creation with a start time would be very useful**

Transfer Matrix - Replica Creation

src \ dst	DESY-DCACHE	SARA-DCACHE	PIC-DCACHE	EULAKE-1	LAPP-DCACHE	IN2P3-CC-DCACHE	CNAF-STORM	ALPAMED-DPM	GSI-ROOT	INFN-NA-DPM	LAPP-WEBDAV	INFN-NA-DPM-FED	INFN-ROMA1
DESY-DCACHE	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA
SARA-DCACHE	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA
PIC-DCACHE	100%	100%	NO DATA	100%	100%	100%	100%	100%	7.1%	100%	NO DATA	NO DATA	NO DATA
EULAKE-1	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA
LAPP-DCACHE	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA
IN2P3-CC-DCACHE	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA
CNAF-STORM	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA
ALPAMED-DPM	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA
GSI-ROOT	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA
INFN-NA-DPM	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA
LAPP-WEBDAV	100%	100%	NO DATA	100%	NO DATA	100%	100%	100%	100%	100%	NO DATA	NO DATA	NO DATA
INFN-NA-DPM-FED	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA
INFN-ROMA1	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA

SKA Tests

Dress Rehearsal

- Pulsar observations
  Simulate two sets of SKA pulsar observations where pulsar search and single pulse search observations are occurring commensally. The data (50MB) is generated at a different cadence and managed via separate scopes.
- QoS lifecycle
  This test attempts to replicate a data lifecycle hourly based on QOS labels on RSEs in the datalake.
- SDC1 data movement
  Move SDC1 data to more performant storage (FAST QoS) to simulate “making data available” to the end user.

Post dress rehearsal

- High throughput test
  Initiate a transfer of a significant volume of data in an attempt to flood the corresponding network link, and understand the achievable throughput between sites.
- Million file test

SDC1 data movement: Description

Test description:

- Dataset used SKA_SKAO_TEAM:SDC1, replications done with a lifetime of 2 weeks
- Test runs with a dataset level rule replicating to FAST QoS. The replicas are made on LAPP-WEBDAV but the rule remains in STUCK state.
- Next run with file level rules to replicate to FAST storage for 2 weeks, not expected to initiate any transfers.
- **Missing file SKA_SKAO_TEAM:560_1000.fits (no existing replicas)**

Point to note

- HIPS versions of the SDC1 data were added to the dataset a week after dress rehearsal and the active dataset level rule for FAST QOS automatically made replicas for these newly added files on LAPP-WEBDAV as well.

SDC1 data movement: Results

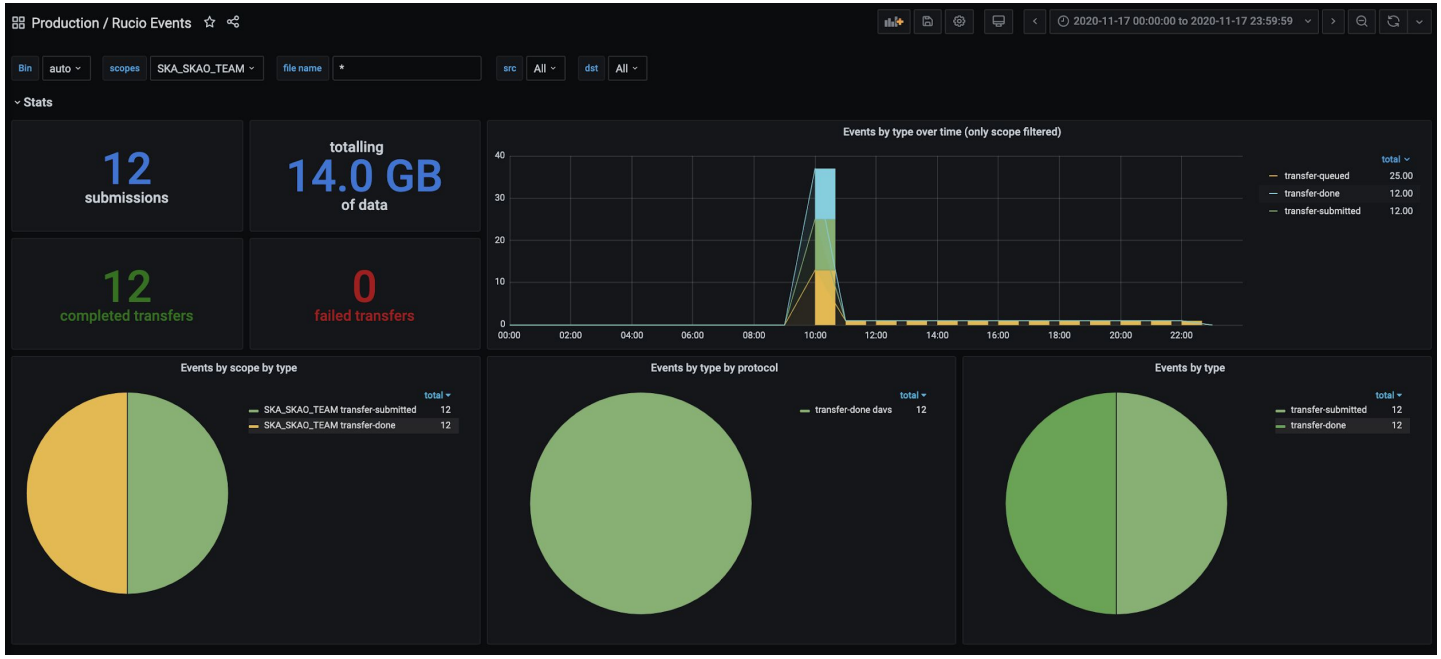

SKA Tests

Dress Rehearsal

- Pulsar observations
  Simulate two sets of SKA pulsar observations where pulsar search and single pulse search observations are occurring commensally. The data (50MB) is generated at a different cadence and managed via separate scopes.
- QoS lifecycle
  This test attempts to replicate a data lifecycle hourly based on QOS labels on RSEs in the datalake.
- SDC1 data movement
  Move SDC1 data to more performant storage (FAST QoS) to simulate “making data available” to the end user.

Post dress rehearsal

- High throughput test
  Initiate a transfer of a significant volume of data in an attempt to flood the corresponding network link, and understand the achievable throughput between sites.
- Million file test

High throughput test: Description

Preparation

- The test data uploaded in the form of 1010 1GB files and grouped in the container SKA_SKAO_JOSHI-testing:1G-files-stress-test.
- Uploaded and replicated to DESY-DCACHE, EULAKE-1, ALPAMED-DPM, SARA-DCACHE

Description

- Test was performed from within the rucio-analysis test framework with the config file here.
- On the day of the test, SKA_SKAO_JOSHI-testing:1G-files-stress-test replicated to the IN2P3-CC-DCACHE RSE with a lifetime of 60 days.
- This took <20 min, thus giving an approximate throughput of **6.7 Gbps**.
- The rule is still in REPLICATING state despite the files seemingly ready at the destination. Since we're calculating the replication_duration based on the change in status of the rule, **maybe a polling frequency somewhere needs to be increased to update the rule state more promptly.**

High throughput test: Results

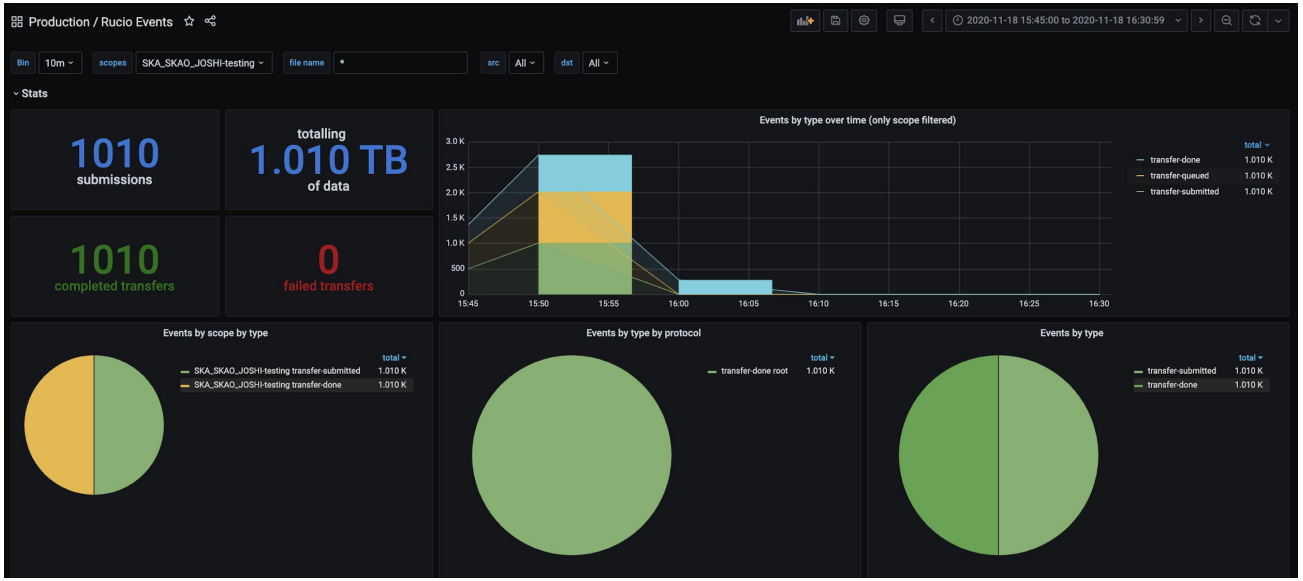

Million File test

Started creating 200-file datasets with 2 week lifetime (so expiring on the 27th) in container SKA_SKAO_TEAM_MFT:million.

Uploading finishes (5295 datasets in the container). Nothing happens, rule remains in INJECT state.

More memory allocated to judge-injector daemon such that it runs on a dedicated node.

Rucio problem fixed, dataset level retention rules created at DESY for expired datasets.

13/11 — 20/11 — 21/11 — 25/11 — 26/11 — 27/11 — 04/12 — 08/12

After initial container level rule creation was timing out, a container level retention rule with async flag set to keep the data at EULAKE-1 for 6 months. Uploading continues to reach 1M files.

Rucio trouble starts. No 'transfers done' seen on dashboard

New Rucio config deployed. Rule moves to replicating state

Rule is in OK state.

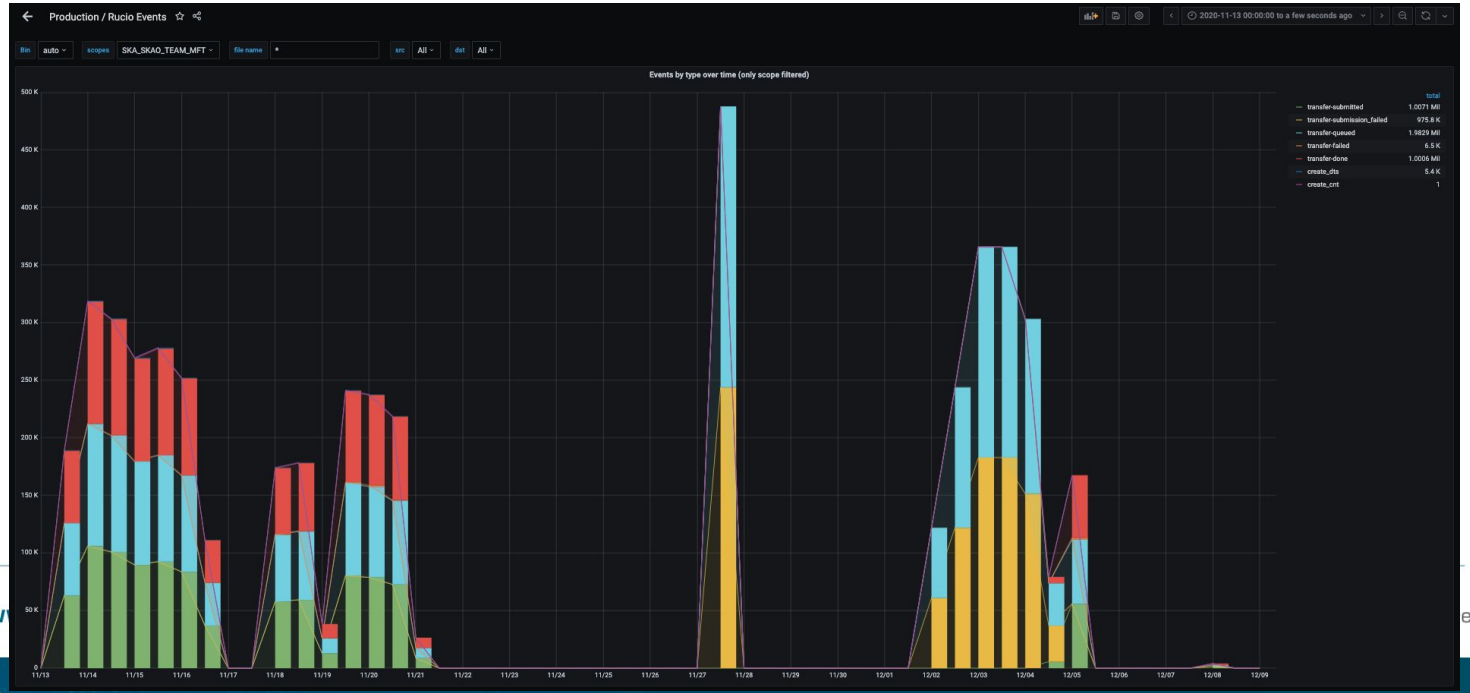

For next time...

- Pulsar observations
  - Simulate data injection possibly via registering data at a non-deterministic RSE
- QoS lifecycle
  - Capture rules with varying expiration dates
- High Throughput
  - Move data to a new destination RSE (TBD) with a source RSE to test a specific link
- Million File Test
  - Get files in the container upto 1 million
  - Move 1M file container to an RSE (TBD: SARA, IN2P3?)

SQUARE KILOMETRE ARRAY

Exploring the Universe with the world's largest radio telescope

Thank you